

# Surrogate Regret Bounds for Bipartite Ranking via Strongly Proper Losses

Shivani Agarwal

Department of Computer Science and Automation  
Indian Institute of Science  
Bangalore 560012, India  
shivani@csa.iisc.ernet.in

July 3, 2012

## Abstract

The problem of bipartite ranking, where instances are labeled positive or negative and the goal is to learn a scoring function that minimizes the probability of mis-ranking a pair of positive and negative instances (or equivalently, that maximizes the area under the ROC curve), has been widely studied in recent years. A dominant theoretical and algorithmic framework for the problem has been to reduce bipartite ranking to pairwise classification; in particular, it is well known that the bipartite ranking regret can be formulated as a pairwise classification regret, which in turn can be upper bounded using usual regret bounds for classification problems. Recently, Kotowski et al. (2011) showed regret bounds for bipartite ranking in terms of the regret associated with balanced versions of the standard (non-pairwise) logistic and exponential losses. In this paper, we show that such (non-pairwise) surrogate regret bounds for bipartite ranking can be obtained in terms of a broad class of proper (composite) losses that we term as *strongly proper*. Our proof technique is much simpler than that of Kotowski et al. (2011), and relies on properties of proper (composite) losses as elucidated recently by Reid and Williamson (2010, 2011) and others. Our result yields explicit surrogate bounds (with no hidden balancing terms) in terms of a variety of strongly proper losses, including for example logistic, exponential, squared and squared hinge losses as special cases. We also obtain tighter surrogate bounds under certain low-noise conditions via a recent result of Cl  men  on and Robbiano (2011).

## 1 Introduction

Ranking problems arise in a variety of applications ranging from information retrieval to recommendation systems and from computational biology to drug discovery, and have been widely studied in machine learning and statistics in the last several years. Recently, there has been much interest in understanding statistical consistency and regret behavior of algorithms for a variety of ranking problems, including various forms of label/subset ranking as well as instance ranking problems [10, 8, 12, 3, 2, 30, 13, 21, 5, 9, 19, 29].

In this paper, we study regret bounds for the bipartite instance ranking problem, where instances are labeled positive or negative and the goal is to learn a scoring function that minimizes the probability of mis-ranking a pair of positive and negative instances, or equivalently, that maximizes the area under the ROC curve [14, 1]. A popular algorithmic and theoretical approach to bipartite

ranking has been to treat the problem as analogous to pairwise classification [17, 18, 14, 20, 7, 8]. Indeed, this approach enjoys theoretical support since the bipartite ranking regret can be formulated as a pairwise classification regret, and therefore any algorithm minimizing the latter over a suitable class of functions will also minimize the ranking regret (this follows formally from results of [8]; see Section 3.1 for a summary). Nevertheless, it has often been observed that algorithms such as AdaBoost, logistic regression, and in some cases even SVMs, which minimize the exponential, logistic, and hinge losses respectively in the standard (non-pairwise) setting, also yield good bipartite ranking performance [11, 20, 24]. For losses such as the exponential or logistic losses, this is not surprising since algorithms minimizing these losses (but not the hinge loss) are known to effectively estimate conditional class probabilities [31]; since the class probability function provides the optimal ranking [8], it is intuitively clear (and follows formally from results in [8, 9]) that any algorithm providing a good approximation to the class probability function should also produce a good ranking. However, there has been very little work so far on quantifying the ranking regret of a scoring function in terms of the regret associated with such surrogate losses.

Recently, [19] showed that the bipartite ranking regret of a scoring function can be upper bounded in terms of the regret associated with *balanced* versions of the standard (non-pairwise) exponential and logistic losses. However their proof technique builds on analyses involving the reduction of bipartite ranking to pairwise classification, and involves analyses specific to the exponential and logistic losses (see Section 3.2). More fundamentally, the balanced losses in their result depend on the underlying distribution and cannot be optimized directly by an algorithm; while it is possible to do so approximately, one then loses the quantitative nature of the bounds.

In this work we obtain quantitative regret bounds for bipartite ranking in terms of a broad class of proper (composite) loss functions that we term *strongly proper*. Our proof technique is considerably simpler than that of [19], and relies on properties of proper (composite) losses as elucidated recently for example in [22, 23, 15, 6]. Our result yields explicit surrogate bounds (with no hidden balancing terms) in terms of a variety of strongly proper (composite) losses, including for example logistic, exponential, squared and squared hinge losses as special cases. We also obtain tighter surrogate bounds under certain low-noise conditions via a recent result of [9].

The paper is organized as follows. In Section 2 we formally set up the bipartite instance ranking problem and definitions related to loss functions and regret, and provide background on proper (composite) losses. Section 3 summarizes related work that provides the background for our study, namely the reduction of bipartite ranking to pairwise binary classification and the result of [19]. In Section 4 we define and characterize strongly proper losses. Section 5 contains our main result, namely a bound on the bipartite ranking regret in terms of the regret associated with any strongly proper loss, together with several examples. Section 6 gives a tighter bound under certain low-noise conditions via a recent result of [9]. We conclude with a brief discussion and some open questions in Section 7.

## 2 Formal Setup, Preliminaries, and Background

This section provides background on the bipartite ranking problem, binary loss functions and regret, and proper (composite) losses.

## 2.1 Bipartite Ranking

As in binary classification, in bipartite ranking there is an instance space  $\mathcal{X}$  and binary labels  $\mathcal{Y} = \{\pm 1\}$ , with an unknown distribution  $D$  on  $\mathcal{X} \times \{\pm 1\}$ . For  $(X, Y) \sim D$  and  $x \in \mathcal{X}$ , we denote  $\eta(x) = \mathbf{P}(Y = 1 \mid X = x)$  and  $p = \mathbf{P}(Y = 1)$ . Given i.i.d. examples  $(X_1, Y_1), \dots, (X_n, Y_n) \sim D$ , the goal is to learn a scoring function  $f : \mathcal{X} \rightarrow \mathbb{R}^*$  (where  $\mathbb{R}^* = [-\infty, \infty]$ ) that assigns higher scores to positive instances than to negative ones.<sup>1</sup> Specifically, the goal is to learn a scoring function  $f$  with low *ranking error* (or *ranking risk*), defined as<sup>2</sup>

$$\text{er}_D^{\text{rank}}[f] = \mathbf{E} \left[ \mathbf{1}((Y - Y')(f(X) - f(X')) < 0) + \frac{1}{2} \mathbf{1}(f(X) = f(X')) \mid Y \neq Y' \right], \quad (1)$$

where  $(X, Y), (X', Y')$  are assumed to be drawn i.i.d. from  $D$ , and  $\mathbf{1}(\cdot)$  is 1 if its argument is true and 0 otherwise; thus the ranking error of  $f$  is simply the probability that a randomly drawn positive instance receives a lower score under  $f$  than a randomly drawn negative instance, with ties broken uniformly at random. The *optimal ranking error* (or *Bayes ranking error* or *Bayes ranking risk*) can be seen to be

$$\text{er}_D^{\text{rank},*} = \inf_{f: \mathcal{X} \rightarrow \mathbb{R}^*} \text{er}_D^{\text{rank}}[f] \quad (2)$$

$$= \frac{1}{2p(1-p)} \mathbf{E}_{X, X'} \left[ \min \left( \eta(X)(1 - \eta(X')), \eta(X')(1 - \eta(X)) \right) \right]. \quad (3)$$

The *ranking regret* of a scoring function  $f : \mathcal{X} \rightarrow \mathbb{R}^*$  is then simply

$$\text{regret}_D^{\text{rank}}[f] = \text{er}_D^{\text{rank}}[f] - \text{er}_D^{\text{rank},*}. \quad (4)$$

We will be interested in upper bounding the ranking regret of a scoring function  $f$  in terms of its regret with respect to certain other (binary) loss functions. In particular, the loss functions we consider will belong to the class of *proper* (composite) loss functions. Below we briefly review some standard notions related to loss functions and regret, and then discuss some properties of proper (composite) losses.

## 2.2 Loss Functions, Regret, and Conditional Risks and Regret

Assume again a probability distribution  $D$  on  $\mathcal{X} \times \{\pm 1\}$  as above. Given a prediction space  $\hat{\mathcal{Y}} \subseteq \mathbb{R}^*$ , a binary *loss function*  $\ell : \{\pm 1\} \times \hat{\mathcal{Y}} \rightarrow \mathbb{R}_+^*$  (where  $\mathbb{R}_+^* = [0, \infty]$ ) assigns a penalty  $\ell(y, \hat{y})$  for predicting  $\hat{y} \in \hat{\mathcal{Y}}$  when the true label is  $y \in \{\pm 1\}$ .<sup>3</sup> For any such loss  $\ell$ , the  $\ell$ -*error* (or  $\ell$ -*risk*) of a function  $f : \mathcal{X} \rightarrow \hat{\mathcal{Y}}$  is defined as

$$\text{er}_D^\ell[f] = \mathbf{E}_{(X, Y) \sim D} [\ell(Y, f(X))], \quad (5)$$

and the *optimal  $\ell$ -error* (or *optimal  $\ell$ -risk* or *Bayes  $\ell$ -risk*) is defined as

$$\text{er}_D^{\ell,*} = \inf_{f: \mathcal{X} \rightarrow \hat{\mathcal{Y}}} \text{er}_D^\ell[f]. \quad (6)$$

<sup>1</sup>Most algorithms learn real-valued functions; we also allow values  $-\infty$  and  $\infty$  for technical reasons.

<sup>2</sup>We assume measurability conditions where necessary.

<sup>3</sup>Most loss functions take values in  $\mathbb{R}_+$ , but some loss functions (such as the logistic loss, described later) can assign a loss of  $\infty$  to certain label-prediction pairs.

The  $\ell$ -regret of a function  $f : \mathcal{X} \rightarrow \widehat{\mathcal{Y}}$  is the difference of its  $\ell$ -error from the optimal  $\ell$ -error:

$$\text{regret}_D^\ell[f] = \text{er}_D^\ell[f] - \text{er}_D^{\ell,*}. \quad (7)$$

The *conditional  $\ell$ -risk*  $L_\ell : [0, 1] \times \widehat{\mathcal{Y}} \rightarrow \mathbb{R}_+^*$  is defined as<sup>4</sup>

$$L_\ell(\eta, \widehat{y}) = \mathbf{E}_{Y \sim \eta}[\ell(Y, \widehat{y})] = \eta \ell(1, \widehat{y}) + (1 - \eta) \ell(-1, \widehat{y}), \quad (8)$$

where  $Y \sim \eta$  denotes a  $\{\pm 1\}$ -valued random variable taking value  $+1$  with probability  $\eta$ . The *conditional Bayes  $\ell$ -risk*  $H_\ell : [0, 1] \rightarrow \mathbb{R}_+^*$  is defined as

$$H_\ell(\eta) = \inf_{\widehat{y} \in \widehat{\mathcal{Y}}} L_\ell(\eta, \widehat{y}). \quad (9)$$

The *conditional  $\ell$ -regret*  $R_\ell : [0, 1] \times \widehat{\mathcal{Y}} \rightarrow \mathbb{R}_+^*$  is then simply

$$R_\ell(\eta, \widehat{y}) = L_\ell(\eta, \widehat{y}) - H_\ell(\eta). \quad (10)$$

Clearly, we have for  $f : \mathcal{X} \rightarrow \widehat{\mathcal{Y}}$ ,

$$\text{er}_D^\ell[f] = \mathbf{E}_X[L_\ell(\eta(X), f(X))], \quad (11)$$

and

$$\text{er}_D^{\ell,*} = \mathbf{E}_X[H_\ell(\eta(X))]. \quad (12)$$

We note the following:

**Lemma 1.** *For any  $\widehat{\mathcal{Y}} \subseteq \mathbb{R}^*$  and binary loss  $\ell : \{\pm 1\} \times \widehat{\mathcal{Y}} \rightarrow \mathbb{R}_+^*$ , the conditional Bayes  $\ell$ -risk  $H_\ell$  is a concave function on  $[0, 1]$ .*

The proof follows simply by observing that  $H_\ell$  is defined as the pointwise infimum of a family of linear (and therefore concave) functions, and therefore is itself concave.

## 2.3 Proper and Proper Composite Losses

In this section we review some background material related to proper and proper composite losses, as studied recently in [22, 23, 15, 6]. While the material is meant to be mostly a review, some of the exposition is simplified compared to previous presentations, and we include a new, simple proof of an important fact (Theorem 4).

**Proper Losses.** We start by considering binary class probability estimation (CPE) loss functions that operate on the prediction space  $\widehat{\mathcal{Y}} = [0, 1]$ . A binary CPE loss function  $c : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$  is said to be *proper* if for all  $\eta \in [0, 1]$ ,

$$\eta \in \arg \min_{\widehat{\eta} \in [0, 1]} L_c(\eta, \widehat{\eta}), \quad (13)$$

and *strictly proper* if the minimizer is unique for all  $\eta \in [0, 1]$ . Equivalently,  $c$  is proper if for all  $\eta \in [0, 1]$ ,  $H_c(\eta) = L_c(\eta, \eta)$ , and strictly proper if  $H_c(\eta) < L_c(\eta, \widehat{\eta})$  for all  $\widehat{\eta} \neq \eta$ . We have the following basic result:

---

<sup>4</sup>Note that we overload notation by using  $\eta$  here to refer to a number in  $[0, 1]$ ; the usage should be clear from context.

**Lemma 2** ([15, 26]). *Let  $c : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$  be a binary CPE loss. If  $c$  is proper, then  $c(1, \cdot)$  is a decreasing function on  $[0, 1]$  and  $c(-1, \cdot)$  is an increasing function. If  $c$  is strictly proper, then  $c(1, \cdot)$  is strictly decreasing on  $[0, 1]$  and  $c(-1, \cdot)$  is strictly increasing.*

We will find it useful to consider *regular* proper losses. As in [15], we say a binary CPE loss  $c : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$  is *regular* if  $c(1, \hat{\eta}) \in \mathbb{R}_+ \forall \hat{\eta} \in (0, 1]$  and  $c(-1, \hat{\eta}) \in \mathbb{R}_+ \forall \hat{\eta} \in [0, 1)$ , i.e. if  $c(y, \hat{\eta})$  is finite for all  $y, \hat{\eta}$  except possibly for  $c(1, 0)$  and  $c(-1, 1)$ , which are allowed to be infinite. The following characterization of regular proper losses is well known (see also [15]):

**Theorem 3** ([25]). *A regular binary CPE loss  $c : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$  is proper if and only if for all  $\eta, \hat{\eta} \in [0, 1]$  there exists a superderivative  $H'_c(\hat{\eta})$  of  $H_c$  at  $\hat{\eta}$  such that<sup>5</sup>*

$$L_c(\eta, \hat{\eta}) = H_c(\hat{\eta}) + (\eta - \hat{\eta}) \cdot H'_c(\hat{\eta}).$$

The following is a characterization of strict properness of a proper loss  $c$  in terms of its conditional Bayes risk  $H_c$ :

**Theorem 4.** *A proper loss  $c : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$  is strictly proper if and only if  $H_c$  is strictly concave.*

This result can be proved in several ways. A proof in [15] is attributed to an argument in [16]. If  $H_c$  is twice differentiable, an alternative proof follows from a result in [6, 26], which shows that a proper loss  $c$  is strictly proper if and only if its ‘weight function’  $w_c = -H'_c$  satisfies  $w_c(\eta) > 0$  for all except at most countably many points  $\eta \in [0, 1]$ ; by a very recent result of [27], this condition is equivalent to strict convexity of the function  $-H_c$ , or equivalently, strict concavity of  $H_c$ . Here we give a third, self-contained proof of the above result that is derived from first principles, and that will be helpful when we study strongly proper losses in Section 4.

*Proof of Theorem 4.* Let  $c : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$  be a proper loss. For the ‘if’ direction, assume  $H_c$  is strictly concave. Let  $\eta, \hat{\eta} \in [0, 1]$  such that  $\hat{\eta} \neq \eta$ . Then we have

$$\begin{aligned} L_c(\eta, \hat{\eta}) - H_c(\eta) &= L_c(\eta, \hat{\eta}) + H_c(\hat{\eta}) - H_c(\hat{\eta}) - H_c(\eta) \\ &= L_c(\eta, \hat{\eta}) + H_c(\hat{\eta}) - 2\left(\frac{1}{2}H_c(\eta) + \frac{1}{2}H_c(\hat{\eta})\right) \\ &> L_c(\eta, \hat{\eta}) + H_c(\hat{\eta}) - 2H_c\left(\frac{\eta + \hat{\eta}}{2}\right) \\ &= 2\left(\left(\frac{\eta + \hat{\eta}}{2}\right)c(1, \hat{\eta}) + \left(1 - \frac{\eta + \hat{\eta}}{2}\right)c(-1, \hat{\eta})\right) - 2H_c\left(\frac{\eta + \hat{\eta}}{2}\right) \\ &= 2\left(L_c\left(\frac{\eta + \hat{\eta}}{2}, \hat{\eta}\right) - H_c\left(\frac{\eta + \hat{\eta}}{2}\right)\right) \\ &\geq 0. \end{aligned}$$

Thus  $c$  is strictly proper.

Conversely, to prove the ‘only if’ direction, assume  $c$  is strictly proper. Let  $\eta_1, \eta_2 \in [0, 1]$  such that  $\eta_1 \neq \eta_2$ , and let  $t \in (0, 1)$ . Then we have

$$\begin{aligned} H_c(t\eta_1 + (1-t)\eta_2) &= L_c(t\eta_1 + (1-t)\eta_2, t\eta_1 + (1-t)\eta_2) \\ &= tL_c(\eta_1, t\eta_1 + (1-t)\eta_2) + (1-t)L_c(\eta_2, t\eta_1 + (1-t)\eta_2) \\ &> tH_c(\eta_1) + (1-t)H_c(\eta_2). \end{aligned}$$

---

<sup>5</sup>Here  $u \in \mathbb{R}$  is a superderivative of  $H_c$  at  $\hat{\eta}$  if for all  $\eta \in [0, 1]$ ,  $H_c(\hat{\eta}) - H_c(\eta) \geq u(\hat{\eta} - \eta)$ .

Thus  $H_c$  is strictly concave.  $\square$

**Proper Composite Losses.** The notion of properness can be extended to binary loss functions operating on prediction spaces  $\widehat{\mathcal{Y}}$  other than  $[0, 1]$  via composition with a *link* function  $\psi : [0, 1] \rightarrow \widehat{\mathcal{Y}}$ . Specifically, for any  $\widehat{\mathcal{Y}} \subseteq \mathbb{R}^*$ , a loss function  $\ell : \{\pm 1\} \times \widehat{\mathcal{Y}} \rightarrow \mathbb{R}_+$  is said to be *proper composite* if it can be written as

$$\ell(y, \widehat{y}) = c(y, \psi^{-1}(\widehat{y})) \quad (14)$$

for some proper loss  $c : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$  and strictly increasing (and therefore invertible) link function  $\psi : [0, 1] \rightarrow \widehat{\mathcal{Y}}$ . Proper composite losses have been studied recently in [22, 23, 6], and include several widely used losses such as squared, squared hinge, logistic, and exponential losses.

It is worth noting that for a proper composite loss  $\ell$  formed from a proper loss  $c$ ,  $H_\ell = H_c$ . Moreover, any property associated with the underlying proper loss  $c$  can also be used to describe the composite loss  $\ell$ ; thus we will refer to a proper composite loss  $\ell$  formed from a regular proper loss  $c$  as *regular proper composite*, a composite loss formed from a strictly proper loss as *strictly proper composite*, etc. In Section 4, we will define and characterize *strongly proper* (composite) losses, which we will use to obtain regret bounds for bipartite ranking.

### 3 Related Work

As noted above, a popular theoretical and algorithmic framework for bipartite ranking has been to reduce the problem to pairwise classification. Below we describe this reduction in the context of our setting and notation, and then review the result of [19] which builds on this pairwise reduction.

#### 3.1 Reduction of Bipartite Ranking to Pairwise Binary Classification

For any distribution  $D$  on  $\mathcal{X} \times \{\pm 1\}$ , consider the distribution  $\widetilde{D}$  on  $(\mathcal{X} \times \mathcal{X}) \times \{\pm 1\}$  defined as follows:

1. Sample  $(X, Y)$  and  $(X', Y')$  i.i.d. from  $D$ ;
2. If  $Y = Y'$ , then go to step 1; else set<sup>6</sup>

$$\widetilde{X} = (X, X'), \quad \widetilde{Y} = \text{sign}(Y - Y')$$

and return  $(\widetilde{X}, \widetilde{Y})$ .

Then it is easy to see that, under  $\widetilde{D}$ ,

$$\mathbf{P}(\widetilde{X} = (x, x')) = \frac{\mathbf{P}(X = x) \mathbf{P}(X' = x') (\eta(x)(1 - \eta(x')) + \eta(x')(1 - \eta(x)))}{2p(1 - p)} \quad (15)$$

$$\widetilde{\eta}((x, x')) = \mathbf{P}(\widetilde{Y} = 1 \mid \widetilde{X} = (x, x')) = \frac{\eta(x)(1 - \eta(x'))}{\eta(x)(1 - \eta(x')) + \eta(x')(1 - \eta(x))} \quad (16)$$

$$\widetilde{p} = \mathbf{P}(\widetilde{Y} = 1) = \frac{1}{2}. \quad (17)$$

---

<sup>6</sup>Throughout the paper,  $\text{sign}(u) = +1$  if  $u > 0$  and  $-1$  otherwise.

Moreover, for the 0-1 loss  $\ell_{0-1} : \{\pm 1\} \times \{\pm 1\} \rightarrow \{0, 1\}$  given by  $\ell_{0-1}(y, \hat{y}) = \mathbf{1}(\hat{y} \neq y)$ , we have the following for any pairwise binary classifier  $h : \mathcal{X} \times \mathcal{X} \rightarrow \{\pm 1\}$ :

$$\text{er}_{\tilde{D}}^{0-1}[h] = \mathbf{E}_{(\tilde{X}, \tilde{Y}) \sim \tilde{D}} [\mathbf{1}(h(\tilde{X}) \neq \tilde{Y})] \quad (18)$$

$$\text{er}_{\tilde{D}}^{0-1,*} = \mathbf{E}_{\tilde{X}} [\min(\tilde{\eta}(\tilde{X}), 1 - \tilde{\eta}(\tilde{X}))] \quad (19)$$

$$\text{regret}_{\tilde{D}}^{0-1}[h] = \text{er}_{\tilde{D}}^{0-1}[h] - \text{er}_{\tilde{D}}^{0-1,*}. \quad (20)$$

Now for any scoring function  $f : \mathcal{X} \rightarrow \mathbb{R}^*$ , define  $f_{\text{diff}} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^*$  as

$$f_{\text{diff}}(x, x') = f(x) - f(x'). \quad (21)$$

Then it is easy to see that:

$$\text{er}_D^{\text{rank}}[f] = \text{er}_{\tilde{D}}^{0-1}[\text{sign} \circ f_{\text{diff}}] \quad (22)$$

$$\text{er}_D^{\text{rank},*} = \text{er}_{\tilde{D}}^{0-1,*}, \quad (23)$$

where  $(g \circ f)(u) = g(f(u))$ . The equality in Eq. (23) follows from the fact that the classifier  $h^*(x, x') = \text{sign}(\eta(x) - \eta(x'))$  achieves the Bayes 0-1 risk, i.e.  $\text{er}_{\tilde{D}}^{0-1}[h^*] = \text{er}_{\tilde{D}}^{0-1,*}$  [8]. Thus

$$\text{regret}_D^{\text{rank}}[f] = \text{regret}_{\tilde{D}}^{0-1}[\text{sign} \circ f_{\text{diff}}], \quad (24)$$

and therefore the ranking regret of a scoring function  $f : \mathcal{X} \rightarrow \mathbb{R}^*$  can be analyzed via upper bounds on the 0-1 regret of the pairwise classifier  $(\text{sign} \circ f_{\text{diff}}) : \mathcal{X} \times \mathcal{X} \rightarrow \{\pm 1\}$ .<sup>7</sup>

In particular, as noted in [8], applying a result of [4], we can upper bound the pairwise 0-1 regret above in terms of the pairwise  $\ell_\phi$ -regret associated with any classification-calibrated margin loss  $\ell_\phi : \{\pm 1\} \times \mathbb{R}^* \rightarrow \mathbb{R}_+^*$ , i.e. any loss of the form  $\ell_\phi(y, \hat{y}) = \phi(y\hat{y})$  for some function  $\phi : \mathbb{R}^* \rightarrow \mathbb{R}_+^*$  satisfying  $\forall \eta \in [0, 1], \eta \neq \frac{1}{2}$ ,<sup>8</sup>

$$\hat{y}^* \in \arg \min_{\hat{y} \in \mathbb{R}^*} L_\phi(\eta, \hat{y}) \implies \hat{y}^*(\eta - \frac{1}{2}) > 0. \quad (25)$$

We note in particular that for every proper composite margin loss, the associated link function  $\psi$  satisfies  $\psi(\frac{1}{2}) = 0$  [22], and therefore every strictly proper composite margin loss is classification-calibrated in the sense above.<sup>9</sup>

**Theorem 5** ([4]; see also [8]). *Let  $\phi : \mathbb{R}^* \rightarrow \mathbb{R}_+^*$  be such that the margin loss  $\ell_\phi : \{\pm 1\} \times \mathbb{R}^* \rightarrow \mathbb{R}_+^*$  defined as  $\ell_\phi(y, \hat{y}) = \phi(y\hat{y})$  is classification-calibrated as above. Then  $\exists$  strictly increasing function  $g_\phi : \mathbb{R}_+^* \rightarrow [0, 1]$  with  $g_\phi(0) = 0$  such that for any  $\tilde{f} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^*$ ,*

$$\text{regret}_{\tilde{D}}^{0-1}[\text{sign} \circ \tilde{f}] \leq g_\phi(\text{regret}_{\tilde{D}}^\phi[\tilde{f}]).$$

<sup>7</sup>Note that the setting here is somewhat different from that of [3] and [2], who consider a *subset* version of bipartite ranking where each instance consists of some finite subset of objects to be ranked; there also the problem is reduced to a (subset) pairwise classification problem, and it is shown that given any (subset) pairwise classifier  $h$ , a subset ranking function  $f$  can be constructed such that the resulting subset ranking regret is at most twice the subset pairwise classification regret of  $h$  [3], or in expectation at most equal to the pairwise classification regret of  $h$  [2].

<sup>8</sup>We abbreviate  $L_\phi = L_{\ell_\phi}$ ,  $\text{er}_D^\phi = \text{er}_D^{\ell_\phi}$ , etc.

<sup>9</sup>We note that in general, every strictly proper (composite) loss is classification-calibrated with respect to any cost-sensitive zero-one loss, using a more general definition of classification calibration with an appropriate threshold (e.g. see [22]).

[4] give a construction for  $g_\phi$ ; in particular, for the exponential loss given by  $\phi_{\text{exp}}(u) = e^{-u}$  and logistic loss given by  $\phi_{\text{log}}(u) = \ln(1 + e^{-u})$ , both of which are strictly proper composite losses (see Section 5.2) and are therefore classification-calibrated, one has

$$g_{\text{exp}}(z) \leq \sqrt{2z} \quad (26)$$

$$g_{\text{log}}(z) \leq \sqrt{2z}. \quad (27)$$

As we describe below, [19] build on these observations to bound the ranking regret in terms of the regret associated with balanced versions of the exponential and logistic losses.

### 3.2 Result of Kotlowski et al. (2011)

For any binary loss  $\ell : \{\pm 1\} \times \hat{\mathcal{Y}} \rightarrow \mathbb{R}_+^*$ , consider defining a *balanced* loss  $\ell_{\text{bal}} : \{\pm 1\} \times \hat{\mathcal{Y}} \rightarrow \mathbb{R}_+^*$  as

$$\ell_{\text{bal}}(y, \hat{y}) = \frac{1}{2p} \ell(1, \hat{y}) \cdot \mathbf{1}(y = 1) + \frac{1}{2(1-p)} \ell(-1, \hat{y}) \cdot \mathbf{1}(y = -1). \quad (28)$$

Note that such a balanced loss depends on the underlying distribution  $D$  via  $p = \mathbf{P}(Y = 1)$ . Then [19] show the following, via analyses specific to the exponential and logistic losses:

**Theorem 6** ([19]). *For any  $f : \mathcal{X} \rightarrow \mathbb{R}^*$ ,*

$$\begin{aligned} \text{regret}_D^{\text{exp}}[f_{\text{diff}}] &\leq \frac{9}{4} \text{regret}_D^{\text{exp, bal}}[f] \\ \text{regret}_D^{\text{log}}[f_{\text{diff}}] &\leq 2 \text{regret}_D^{\text{log, bal}}[f]. \end{aligned}$$

Combining this with the results of Eq. (24), Theorem 5, and Eqs. (26-27) then gives the following bounds on the ranking regret of any scoring function  $f : \mathcal{X} \rightarrow \mathbb{R}^*$  in terms of the (non-pairwise) balanced exponential and logistic regrets of  $f$ :

$$\text{regret}_D^{\text{rank}}[f] \leq \frac{3}{\sqrt{2}} \sqrt{\text{regret}_D^{\text{exp, bal}}[f]} \quad (29)$$

$$\text{regret}_D^{\text{rank}}[f] \leq 2 \sqrt{\text{regret}_D^{\text{log, bal}}[f]}. \quad (30)$$

This suggests that an algorithm that produces a function  $f : \mathcal{X} \rightarrow \mathbb{R}^*$  with low balanced exponential or logistic regret will also have low ranking regret. Unfortunately, since the balanced losses depend on the unknown distribution  $D$ , they cannot be optimized by an algorithm directly.<sup>10</sup> [19] provide some justification for why in certain situations, minimizing the usual exponential or logistic loss may also minimize the balanced versions of these losses; however, by doing so, one loses the quantitative nature of the above bounds. Below we obtain upper bounds on the ranking regret of a function  $f$  directly in terms of its loss-based regret (with no balancing terms) for a wide range of proper (composite) loss functions that we term *strongly proper*, including the exponential and logistic losses as special cases.

---

<sup>10</sup>We note it is possible to optimize approximately balanced losses, e.g. by estimating  $p$  from the data.



## 4 Strongly Proper Losses

We define strongly proper losses as follows:

**Definition 7.** Let  $c : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$  be a binary CPE loss and let  $\lambda > 0$ . We say  $c$  is  $\lambda$ -strongly proper if for all  $\eta, \hat{\eta} \in [0, 1]$ ,

$$L_c(\eta, \hat{\eta}) - H_c(\eta) \geq \frac{\lambda}{2}(\eta - \hat{\eta})^2.$$

We have the following necessary and sufficient conditions for strong properness:

**Lemma 8.** Let  $\lambda > 0$ . If  $c : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$  is  $\lambda$ -strongly proper, then  $H_c$  is  $\lambda$ -strongly concave.

*Proof.* The proof is similar to the ‘only if’ direction in the proof of Theorem 4. Let  $c$  be  $\lambda$ -strongly proper. Let  $\eta_1, \eta_2 \in [0, 1]$  such that  $\eta_1 \neq \eta_2$ , and let  $t \in (0, 1)$ . Then we have

$$\begin{aligned} H_c(t\eta_1 + (1-t)\eta_2) &= L_c(t\eta_1 + (1-t)\eta_2, t\eta_1 + (1-t)\eta_2) \\ &= t L_c(\eta_1, t\eta_1 + (1-t)\eta_2) + (1-t) L_c(\eta_2, t\eta_1 + (1-t)\eta_2) \\ &\geq t \left( H_c(\eta_1) + \frac{\lambda}{2}(1-t)^2(\eta_1 - \eta_2)^2 \right) + (1-t) \left( H_c(\eta_2) + \frac{\lambda}{2}t^2(\eta_1 - \eta_2)^2 \right) \\ &= t H_c(\eta_1) + (1-t) H_c(\eta_2) + \frac{\lambda}{2}t(1-t)(\eta_1 - \eta_2)^2. \end{aligned}$$

Thus  $H_c$  is  $\lambda$ -strongly concave. □

**Lemma 9.** Let  $\lambda > 0$  and let  $c : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$  be a regular proper loss. If  $H_c$  is  $\lambda$ -strongly concave, then  $c$  is  $\lambda$ -strongly proper.

*Proof.* Let  $\eta, \hat{\eta} \in [0, 1]$ . By Theorem 3, there exists a superderivative  $H'_c(\hat{\eta})$  of  $H_c$  at  $\hat{\eta}$  such that

$$L_c(\eta, \hat{\eta}) = H_c(\hat{\eta}) + (\eta - \hat{\eta}) \cdot H'_c(\hat{\eta}).$$

This gives

$$\begin{aligned} L_c(\eta, \hat{\eta}) - H_c(\eta) &= H_c(\hat{\eta}) - H_c(\eta) + (\eta - \hat{\eta}) \cdot H'_c(\hat{\eta}) \\ &\geq \frac{\lambda}{2}(\hat{\eta} - \eta)^2, \quad \text{since } H_c \text{ is } \lambda\text{-strongly concave.} \end{aligned}$$

Thus  $c$  is  $\lambda$ -strongly proper. □

This gives us the following characterization of strong properness for regular proper losses:

**Theorem 10.** Let  $\lambda > 0$  and let  $c : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$  be a regular proper loss. Then  $c$  is  $\lambda$ -strongly proper if and only if  $H_c$  is  $\lambda$ -strongly concave.

Several examples of strongly proper (composite) losses will be provided in Section 5.2 and Section 5.3. Theorem 10 will form our main tool in establishing strong properness of many of these loss functions.

## 5 Regret Bounds via Strongly Proper Losses

We start by recalling the following result of [8] (adapted to account for ties, and for the conditioning on  $Y \neq Y'$ ):

**Theorem 11** ([8]). *For any  $f : \mathcal{X} \rightarrow \mathbb{R}^*$ ,*

$$\begin{aligned} \text{regret}_D^{\text{rank}}[f] &= \frac{1}{2p(1-p)} \mathbf{E}_{X, X'} \left[ |\eta(X) - \eta(X')| \cdot \left( \mathbf{1}((f(X) - f(X'))(\eta(X) - \eta(X')) < 0) \right. \right. \\ &\quad \left. \left. + \frac{1}{2} \mathbf{1}(f(X) = f(X')) \right) \right]. \end{aligned}$$

As noted by [9], this leads to the following corollary on the regret of any plug-in ranking function based on an estimate  $\hat{\eta}$ :

**Corollary 12.** *For any  $\hat{\eta} : \mathcal{X} \rightarrow [0, 1]$ ,*

$$\text{regret}_D^{\text{rank}}[\hat{\eta}] \leq \frac{1}{p(1-p)} \mathbf{E}_X [|\hat{\eta}(X) - \eta(X)|].$$

For completeness, a proof is given in Appendix A. We are now ready to prove our main result.

### 5.1 Main Result

**Theorem 13.** *Let  $\hat{\mathcal{Y}} \subseteq \mathbb{R}^*$  and let  $\lambda > 0$ . Let  $\ell : \{\pm 1\} \times \hat{\mathcal{Y}} \rightarrow \mathbb{R}_+^*$  be a  $\lambda$ -strongly proper composite loss. Then for any  $f : \mathcal{X} \rightarrow \hat{\mathcal{Y}}$ ,*

$$\text{regret}_D^{\text{rank}}[f] \leq \frac{\sqrt{2}}{p(1-p)\sqrt{\lambda}} \sqrt{\text{regret}_D^\ell[f]}.$$

*Proof.* Let  $c : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$  be a  $\lambda$ -strongly proper loss and  $\psi : [0, 1] \rightarrow \hat{\mathcal{Y}}$  be a (strictly increasing) link function such that  $\ell(y, \hat{y}) = c(y, \psi^{-1}(\hat{y}))$  for all  $y \in \{\pm 1\}, \hat{y} \in \hat{\mathcal{Y}}$ . Let  $f : \mathcal{X} \rightarrow \hat{\mathcal{Y}}$ . Then we have,

$$\begin{aligned} \text{regret}_D^{\text{rank}}[f] &= \text{regret}_D^{\text{rank}}[\psi^{-1} \circ f], \quad \text{since } \psi \text{ is strictly increasing} \\ &\leq \frac{1}{p(1-p)} \mathbf{E}_X [|\psi^{-1}(f(X)) - \eta(X)|], \quad \text{by Corollary 12} \\ &= \frac{1}{p(1-p)} \sqrt{\left( \mathbf{E}_X [|\psi^{-1}(f(X)) - \eta(X)|] \right)^2} \\ &\leq \frac{1}{p(1-p)} \sqrt{\mathbf{E}_X \left[ (\psi^{-1}(f(X)) - \eta(X))^2 \right]}, \\ &\quad \text{by convexity of } \phi(u) = u^2 \text{ and Jensen's inequality} \\ &\leq \frac{1}{p(1-p)} \sqrt{\frac{2}{\lambda} \mathbf{E}_X [R_c(\eta(X), \psi^{-1}(f(X)))]}, \quad \text{since } c \text{ is } \lambda\text{-strongly proper} \\ &= \frac{1}{p(1-p)} \sqrt{\frac{2}{\lambda} \mathbf{E}_X [R_\ell(\eta(X), (f(X)))]} \\ &= \frac{\sqrt{2}}{p(1-p)\sqrt{\lambda}} \sqrt{\text{regret}_D^\ell[f]}. \end{aligned}$$

□

Theorem 13 shows that for any strongly proper composite loss  $\ell : \{\pm 1\} \times \widehat{\mathcal{Y}} \rightarrow \mathbb{R}_+^*$ , a function  $f : \mathcal{X} \rightarrow \widehat{\mathcal{Y}}$  with low  $\ell$ -regret will also have low ranking regret. Below we give several examples of such strongly proper (composite) loss functions; properties of some of these losses are summarized in Table 1.

## 5.2 Examples

**Example 1** (Exponential loss). *The exponential loss  $\ell_{\text{exp}} : \{\pm 1\} \times \mathbb{R}^* \rightarrow \mathbb{R}_+^*$  defined as*

$$\ell_{\text{exp}}(y, \hat{y}) = e^{-y\hat{y}}$$

*is a proper composite loss with associated proper loss  $c_{\text{exp}} : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$  and link function  $\psi_{\text{exp}} : [0, 1] \rightarrow \mathbb{R}^*$  given by*

$$c_{\text{exp}}(y, \hat{\eta}) = \left( \frac{1 - \hat{\eta}}{\hat{\eta}} \right)^{y/2}; \quad \psi_{\text{exp}}(\hat{\eta}) = \frac{1}{2} \ln \left( \frac{\hat{\eta}}{1 - \hat{\eta}} \right).$$

*It is easily verified that  $c_{\text{exp}}$  is regular. Moreover, it can be seen that*

$$H_{\text{exp}}(\eta) = 2\sqrt{\eta(1 - \eta)},$$

*with*

$$-H_{\text{exp}}''(\eta) = \frac{1}{2(\eta(1 - \eta))^{3/2}} \geq 4 \quad \forall \eta \in [0, 1].$$

*Thus  $H_{\text{exp}}$  is 4-strongly concave, and so by Theorem 10, we have  $\ell_{\text{exp}}$  is 4-strongly proper composite. Therefore applying Theorem 13 we have for any  $f : \mathcal{X} \rightarrow \mathbb{R}^*$ ,*

$$\text{regret}_D^{\text{rank}}[f] \leq \frac{1}{\sqrt{2}p(1 - p)} \sqrt{\text{regret}_D^{\text{exp}}[f]}.$$

**Example 2** (Logistic loss). *The logistic loss  $\ell_{\text{log}} : \{\pm 1\} \times \mathbb{R}^* \rightarrow \mathbb{R}_+^*$  defined as*

$$\ell_{\text{log}}(y, \hat{y}) = \ln(1 + e^{-y\hat{y}})$$

*is a proper composite loss with associated proper loss  $c_{\text{log}} : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$  and link function  $\psi_{\text{log}} : [0, 1] \rightarrow \mathbb{R}^*$  given by*

$$c_{\text{log}}(1, \hat{\eta}) = -\ln \hat{\eta}; \quad c_{\text{log}}(-1, \hat{\eta}) = -\ln(1 - \hat{\eta}); \quad \psi_{\text{log}}(\hat{\eta}) = \ln \left( \frac{\hat{\eta}}{1 - \hat{\eta}} \right).$$

*Again, it is easily verified that  $c_{\text{log}}$  is regular. Moreover, it can be seen that*

$$H_{\text{log}}(\eta) = -\eta \ln \eta - (1 - \eta) \ln(1 - \eta),$$

*with*

$$-H_{\text{log}}''(\eta) = \frac{1}{\eta(1 - \eta)} \geq 4 \quad \forall \eta \in [0, 1].$$

*Thus  $H_{\text{log}}$  is 4-strongly concave, and so by Theorem 10, we have  $\ell_{\text{log}}$  is 4-strongly proper composite. Therefore applying Theorem 13 we have for any  $f : \mathcal{X} \rightarrow \mathbb{R}^*$ ,*

$$\text{regret}_D^{\text{rank}}[f] \leq \frac{1}{\sqrt{2}p(1 - p)} \sqrt{\text{regret}_D^{\text{log}}[f]}.$$

**Example 3** (Squared and squared hinge losses). *The (binary) squared loss  $(1 - y\hat{y})^2$  and squared hinge loss  $((1 - y\hat{y})_+)^2$  (where  $u_+ = \max(u, 0)$ ) are generally defined for  $\hat{y} \in \mathbb{R}$ . To obtain class probability estimates from a predicted value  $\hat{y} \in \mathbb{R}$ , one then truncates  $\hat{y}$  to  $[-1, 1]$ , and uses  $\hat{\eta} = \frac{\hat{y}+1}{2}$  [31]. To obtain a proper loss, we can take  $\hat{\mathcal{Y}} = [-1, 1]$ ; in this range, both losses coincide, and we can define  $\ell_{\text{sq}} : \{\pm 1\} \times [-1, 1] \rightarrow [0, 4]$  as*

$$\ell_{\text{sq}}(y, \hat{y}) = (1 - y\hat{y})^2.$$

*This is a proper composite loss with associated proper loss  $c_{\text{sq}} : \{\pm 1\} \times [-1, 1] \rightarrow [0, 4]$  and link function  $\psi_{\text{sq}} : [0, 1] \rightarrow [-1, 1]$  given by*

$$c_{\text{sq}}(1, \hat{\eta}) = 4(1 - \hat{\eta})^2; \quad c_{\text{sq}}(-1, \hat{\eta}) = 4\hat{\eta}^2; \quad \psi_{\text{sq}}(\hat{\eta}) = 2\hat{\eta} - 1.$$

*It can be seen that*

$$L_{\text{sq}}(\eta, \hat{\eta}) = 4\eta(1 - \hat{\eta})^2 + 4(1 - \eta)\hat{\eta}^2$$

*and*

$$H_{\text{sq}}(\eta) = 4\eta(1 - \eta),$$

*so that*

$$L_{\text{sq}}(\eta, \hat{\eta}) - H_{\text{sq}}(\eta) = 4(\eta - \hat{\eta})^2.$$

*Thus  $\ell_{\text{sq}}$  is 8-strongly proper composite, and so applying Theorem 13 we have for any  $f : \mathcal{X} \rightarrow [-1, 1]$ ,*

$$\text{regret}_D^{\text{rank}}[f] \leq \frac{1}{2p(1-p)} \sqrt{\text{regret}_D^{\text{sq}}[f]}.$$

*Note that, if a function  $f : \mathcal{X} \rightarrow \mathbb{R}$  is learned, then our bound in terms of  $\ell_{\text{sq}}$ -regret applies to the ranking regret of an appropriately transformed function  $\bar{f} : \mathcal{X} \rightarrow [-1, 1]$ , such as that obtained by truncating values  $f(x) \notin [-1, 1]$  to the appropriate endpoint  $-1$  or  $1$ :*

$$\bar{f}(x) = \begin{cases} -1 & \text{if } f(x) < -1 \\ f(x) & \text{if } f(x) \in [-1, 1] \\ 1 & \text{if } f(x) > 1. \end{cases}$$

### 5.3 Constructing Strongly Proper Losses

In general, given any concave function  $H : [0, 1] \rightarrow \mathbb{R}_+$ , one can construct a proper loss  $c : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+$  with  $H_c = H$  as follows:

$$c(1, \hat{\eta}) = H(\hat{\eta}) + (1 - \hat{\eta})H'(\hat{\eta}) \tag{31}$$

$$c(-1, \hat{\eta}) = H(\hat{\eta}) - \hat{\eta}H'(\hat{\eta}), \tag{32}$$

where  $H'(\hat{\eta})$  denotes any superderivative of  $H$  at  $\hat{\eta}$ . It can be verified that this gives  $L_c(\eta, \hat{\eta}) = H(\hat{\eta}) + (\eta - \hat{\eta})H'(\hat{\eta})$  for all  $\eta, \hat{\eta} \in [0, 1]$ , and therefore  $H_c(\eta) = H(\eta)$  for all  $\eta \in [0, 1]$ . Moreover, if  $H$  is such that  $H(\hat{\eta}) + (1 - \hat{\eta})H'(\hat{\eta}) \in \mathbb{R}_+ \forall \hat{\eta} \in (0, 1]$  and  $H(\hat{\eta}) - \hat{\eta}H'(\hat{\eta}) \in \mathbb{R}_+ \forall \hat{\eta} \in [0, 1)$ , then the loss  $c$  constructed above is also regular. Thus, starting with any  $\lambda$ -strongly concave function  $H : [0, 1] \rightarrow \mathbb{R}_+$  satisfying these regularity conditions, any proper composite loss  $\ell$  formed from the loss function  $c$  constructed according to Eqs. (31-32) (and any link function  $\psi$ ) is  $\lambda$ -strongly proper composite.

Table 1: Examples of strongly proper composite losses  $\ell : \{\pm 1\} \times \hat{\mathcal{Y}} \rightarrow \mathbb{R}_+^*$  satisfying the conditions of Theorem 13, together with prediction space  $\hat{\mathcal{Y}}$ , proper loss  $c : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$ , link function  $\psi : [0, 1] \rightarrow \hat{\mathcal{Y}}$ , and strong properness parameter  $\lambda$ .

Loss	$\hat{\mathcal{Y}}$	$\ell(y, \hat{y})$	$c(y, \hat{\eta})$		$\psi(\hat{\eta})$	$\lambda$
			$y = 1$	$y = -1$		
Exponential	$\mathbb{R}^*$	$e^{-y\hat{y}}$	$\sqrt{\frac{1-\hat{\eta}}{\hat{\eta}}}$	$\sqrt{\frac{\hat{\eta}}{1-\hat{\eta}}}$	$\frac{1}{2} \ln \left( \frac{\hat{\eta}}{1-\hat{\eta}} \right)$	4
Logistic	$\mathbb{R}^*$	$\ln(1 + e^{-y\hat{y}})$	$-\ln \hat{\eta}$	$-\ln(1 - \hat{\eta})$	$\ln \left( \frac{\hat{\eta}}{1-\hat{\eta}} \right)$	4
Squared	$[-1, 1]$	$(1 - y\hat{y})^2$	$4(1 - \hat{\eta})^2$	$4\hat{\eta}^2$	$2\hat{\eta} - 1$	8
Spherical	$[0, 1]$	$c(y, \hat{y})$	$1 - \frac{\hat{\eta}}{\sqrt{\hat{\eta}^2 + (1-\hat{\eta})^2}}$	$1 - \frac{1-\hat{\eta}}{\sqrt{\hat{\eta}^2 + (1-\hat{\eta})^2}}$	$\hat{\eta}$	1
Canonical ‘exponential’	$\mathbb{R}^*$	$\sqrt{1 + \left(\frac{\hat{y}}{2}\right)^2} - \frac{y\hat{y}}{2}$	$\sqrt{\frac{1-\hat{\eta}}{\hat{\eta}}}$	$\sqrt{\frac{\hat{\eta}}{1-\hat{\eta}}}$	$\frac{2\hat{\eta}-1}{\sqrt{\hat{\eta}(1-\hat{\eta})}}$	4
Canonical squared	$[-1, 1]$	$\frac{1}{4}(1 - y\hat{y})^2$	$(1 - \hat{\eta})^2$	$\hat{\eta}^2$	$2\hat{\eta} - 1$	2
Canonical spherical	$[-1, 1]$	$1 - \frac{1}{2}(\sqrt{2 - \hat{y}^2} + y\hat{y})$	$1 - \frac{\hat{\eta}}{\sqrt{\hat{\eta}^2 + (1-\hat{\eta})^2}}$	$1 - \frac{1-\hat{\eta}}{\sqrt{\hat{\eta}^2 + (1-\hat{\eta})^2}}$	$\frac{2\hat{\eta}-1}{\sqrt{\hat{\eta}^2 + (1-\hat{\eta})^2}}$	1

**Example 4** (Spherical loss). Consider starting with the function  $H_{\text{spher}} : [0, 1] \rightarrow \mathbb{R}$  defined as

$$H_{\text{spher}}(\eta) = 1 - \sqrt{\eta^2 + (1 - \eta)^2}.$$

Then

$$H'_{\text{spher}}(\eta) = \frac{-(2\eta - 1)}{\sqrt{\eta^2 + (1 - \eta)^2}}$$

and

$$-H''_{\text{spher}}(\eta) = \frac{1}{(\eta^2 + (1 - \eta)^2)^{3/2}} \geq 1 \quad \forall \eta \in [0, 1],$$

and therefore  $H_{\text{spher}}$  is 1-strongly concave. Moreover, since  $H_{\text{spher}}$  and  $H'_{\text{spher}}$  are both bounded, the conditions for regularity are also satisfied. Thus we can use Eqs. (31-32) to construct a 1-strongly proper loss  $c_{\text{spher}} : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}$  as follows:

$$\begin{aligned} c_{\text{spher}}(1, \hat{\eta}) &= H_{\text{spher}}(\hat{\eta}) + (1 - \hat{\eta})H'_{\text{spher}}(\hat{\eta}) = 1 - \frac{\hat{\eta}}{\sqrt{\hat{\eta}^2 + (1 - \hat{\eta})^2}} \\ c_{\text{spher}}(-1, \hat{\eta}) &= H_{\text{spher}}(\hat{\eta}) - \hat{\eta}H'_{\text{spher}}(\hat{\eta}) = 1 - \frac{1 - \hat{\eta}}{\sqrt{\hat{\eta}^2 + (1 - \hat{\eta})^2}}. \end{aligned}$$

Therefore by Theorem 13, we have for any  $f : \mathcal{X} \rightarrow [0, 1]$ ,

$$\text{regret}_D^{\text{rank}}[f] \leq \frac{\sqrt{2}}{p(1-p)} \sqrt{\text{regret}_D^{\text{spher}}[f]}.$$

The loss  $c_{\text{spher}}$  above corresponds to the spherical scoring rule described in [15].

We also note that, for every strictly proper loss  $c : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$ , there is an associated ‘canonical’ link function  $\psi : [0, 1] \rightarrow \hat{\mathcal{Y}}$  defined as

$$\psi(\hat{\eta}) = c(-1, \hat{\eta}) - c(1, \hat{\eta}), \quad (33)$$

where  $\hat{\mathcal{Y}} = \{\psi(\hat{\eta}) : \hat{\eta} \in [0, 1]\}$ . We refer to composite losses comprised of such a strictly proper loss  $c$  with the corresponding canonical link  $\psi$  as *canonical proper composite losses*. Clearly, multiplying  $c$  by a factor  $\alpha > 0$  results in the corresponding canonical link  $\psi$  also being multiplied by  $\alpha$ ; adding a constant (or a function  $\theta(y, \hat{\eta}) = \theta(\hat{\eta})$  that is independent of  $y$ ) to  $c$  has no effect on  $\psi$ . Conversely, given any  $\hat{\mathcal{Y}} \subseteq \mathbb{R}^*$  and any (strictly increasing) link function  $\psi : [0, 1] \rightarrow \hat{\mathcal{Y}}$ , there is a unique strictly proper loss  $c : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$  (up to addition of constants or functions of the form  $\theta(y, \hat{\eta}) = \theta(\hat{\eta})$ ) for which  $\psi$  is canonical; this is obtained using Eqs. (31-32) with  $H$  satisfying  $H'(\hat{\eta}) = -\psi(\hat{\eta})$  (with possible addition of a term  $\theta(\hat{\eta})$  to both  $c(1, \hat{\eta})$  and  $c(-1, \hat{\eta})$  thus constructed). Canonical proper composite losses  $\ell(y, \hat{y})$  have some desirable properties, including for example convexity in their second argument  $\hat{y}$  for each  $y \in \{\pm 1\}$ ; we refer the reader to [6, 22] for further discussion of such properties.

We note that the logistic loss in Example 2 is a canonical proper composite loss. On the other hand, as noted in [6], the link  $\psi_{\text{exp}}$  associated with the exponential loss in Example 1 is not the canonical link for the proper loss  $c_{\text{exp}}$  (see Example 5). The squared loss in Example 3 is almost canonical, modulo a scaling factor; one needs to scale either the link function or the loss appropriately (Example 6).

**Example 5** (Canonical proper composite loss associated with  $c_{\text{exp}}$ ). *Let  $c_{\text{exp}} : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$  be as in Example 1. The corresponding canonical link  $\psi_{\text{exp,can}} : [0, 1] \rightarrow \mathbb{R}^*$  is given by*

$$\psi_{\text{exp,can}}(\hat{\eta}) = \sqrt{\frac{\hat{\eta}}{1-\hat{\eta}}} - \sqrt{\frac{1-\hat{\eta}}{\hat{\eta}}} = \frac{2\hat{\eta}-1}{\sqrt{\hat{\eta}(1-\hat{\eta})}}.$$

*With a little algebra, it can be seen that the resulting canonical proper composite loss  $\ell_{\text{exp,can}} : \{\pm 1\} \times \mathbb{R}^* \rightarrow \mathbb{R}_+^*$  is given by*

$$\ell_{\text{exp,can}}(y, \hat{y}) = \sqrt{1 + \left(\frac{\hat{y}}{2}\right)^2} - \frac{y\hat{y}}{2}.$$

*Since we saw  $c_{\text{exp}}$  is 4-strongly proper, we have  $\ell_{\text{exp,can}}$  is 4-strongly proper composite, and therefore we have from Theorem 13 that for any  $f : \mathcal{X} \rightarrow \mathbb{R}^*$ ,*

$$\text{regret}_D^{\text{rank}}[f] \leq \frac{1}{\sqrt{2}p(1-p)} \sqrt{\text{regret}_D^{\text{exp,can}}[f]}.$$

**Example 6** (Canonical squared loss). *For  $c_{\text{sq}} : \{\pm 1\} \times [0, 1] \rightarrow [0, 4]$  defined as in Example 3, the canonical link  $\psi_{\text{sq,can}} : [0, 1] \rightarrow \hat{\mathcal{Y}}$  is given by*

$$\psi_{\text{sq,can}}(\hat{\eta}) = 4\hat{\eta}^2 - 4(1-\hat{\eta})^2 = 4(2\hat{\eta}-1),$$

*with  $\hat{\mathcal{Y}} = [-4, 4]$ , and the resulting canonical squared loss  $\ell_{\text{sq,can}} : \{\pm 1\} \times [-4, 4] \rightarrow [0, 4]$  is given by*

$$\ell_{\text{sq,can}}(y, \hat{y}) = \left(1 - \frac{y\hat{y}}{4}\right)^2.$$

Since we saw  $c_{\text{sq}}$  is 4-strongly proper, we have  $\ell_{\text{sq},\text{can}}$  is 4-strongly proper composite, giving for any  $f : \mathcal{X} \rightarrow [-4, 4]$ ,

$$\text{regret}_D^{\text{rank}}[f] \leq \frac{1}{2p(1-p)} \sqrt{\text{regret}_D^{\text{sq},\text{can}}[f]}.$$

For practical purposes, this is equivalent to using the loss  $\ell_{\text{sq}} : \{\pm 1\} \times [-1, 1] \rightarrow [0, 4]$  defined in Example 3. Alternatively, we can start with a scaled version of the squared proper loss  $c_{\text{sq}'} : \{\pm 1\} \times [0, 1] \rightarrow [0, 1]$  defined as

$$c_{\text{sq}'}(1, \hat{\eta}) = (1 - \hat{\eta})^2; \quad c_{\text{sq}'}(-1, \hat{\eta}) = \hat{\eta}^2,$$

for which the associated canonical link  $\psi_{\text{sq}',\text{can}} : [0, 1] \rightarrow \hat{\mathcal{Y}}$  is given by

$$\psi_{\text{sq}',\text{can}}(\hat{\eta}) = \hat{\eta}^2 - (1 - \hat{\eta})^2 = 2\hat{\eta} - 1,$$

with  $\hat{\mathcal{Y}} = [-1, 1]$ , and the resulting canonical squared loss  $\ell_{\text{sq}',\text{can}} : \{\pm 1\} \times [-1, 1] \rightarrow [0, 1]$  is given by

$$\ell_{\text{sq}',\text{can}}(y, \hat{y}) = \frac{(1 - y\hat{y})^2}{4}.$$

Again, it can be verified that  $c_{\text{sq}'}$  is regular; in this case  $H_{\text{sq}'}(\eta) = \eta(1 - \eta)$  which is 2-strongly concave, giving that  $H_{\text{sq}',\text{can}}$  is 4-strongly proper composite. Therefore applying Theorem 13 we have for any  $f : \mathcal{X} \rightarrow [-1, 1]$ ,

$$\text{regret}_D^{\text{rank}}[f] \leq \frac{1}{p(1-p)} \sqrt{\text{regret}_D^{\text{sq}',\text{can}}[f]}.$$

Again, for practical purposes, this is equivalent to using the loss  $\ell_{\text{sq}}$  defined in Example 3.

**Example 7** (Canonical spherical loss). For  $c_{\text{spher}} : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}$  defined as in Example 4, the canonical link  $\psi_{\text{spher},\text{can}} : [0, 1] \rightarrow \hat{\mathcal{Y}}$  is given by

$$\psi_{\text{spher},\text{can}}(\hat{\eta}) = \frac{2\hat{\eta} - 1}{\sqrt{\hat{\eta}^2 + (1 - \hat{\eta})^2}},$$

with  $\hat{\mathcal{Y}} = [-1, 1]$ . The resulting canonical spherical loss  $\ell_{\text{spher},\text{can}} : \{\pm 1\} \times [-1, 1] \rightarrow \mathbb{R}$  is given by

$$\ell_{\text{spher},\text{can}}(y, \hat{y}) = 1 - \frac{1}{2} \left( \sqrt{2 - \hat{y}^2} + y\hat{y} \right).$$

Since we saw  $c_{\text{spher}}$  is 1-strongly proper, we have  $\ell_{\text{spher},\text{can}}$  is 1-strongly proper composite, and therefore we have from Theorem 13 that for any  $f : \mathcal{X} \rightarrow [-1, 1]$ ,

$$\text{regret}_D^{\text{rank}}[f] \leq \frac{\sqrt{2}}{p(1-p)} \sqrt{\text{regret}_D^{\text{spher},\text{can}}[f]}.$$

## 6 Tighter Bounds under Low-Noise Conditions

In essence, our results exploit the fact that for a  $\lambda$ -strongly proper composite loss  $\ell$  formed from a  $\lambda$ -strongly proper loss  $c$  and link function  $\psi$ , given any scoring function  $f$ , the  $L_2(\mu)$  distance (where  $\mu$  denotes the marginal density of  $D$  on  $\mathcal{X}$ ) between  $\psi^{-1}(f(X))$  and  $\eta(X)$  (and therefore the  $L_1(\mu)$  distance between  $\psi^{-1}(f(X))$  and  $\eta(X)$ , which gives an upper bound on the ranking risk of  $f$ ) can be upper bounded precisely in terms of the  $\ell$ -regret of  $f$ . From this perspective,  $\hat{\eta} = \psi^{-1} \circ f$  can be treated as a ‘plug-in’ scoring function, which we analyzed via Corollary 12.

Recently, [9] showed that, under certain low-noise assumptions, one can obtain tighter bounds on the ranking risk of a plug-in scoring function  $\hat{\eta} : \mathcal{X} \rightarrow [0, 1]$  than that offered by Corollary 12. Specifically, [9] consider the following noise assumption for bipartite ranking (inspired by the noise condition studied in [28] for binary classification):

**Noise Assumption NA( $\alpha$ )** ( $\alpha \in [0, 1]$ ): *A distribution  $D$  on  $\mathcal{X} \times \{\pm 1\}$  satisfies assumption NA( $\alpha$ ) if  $\exists$  a constant  $C > 0$  such that for all  $x \in \mathcal{X}$  and  $t \in [0, 1]$ ,*

$$\mathbf{P}_X(|\eta(X) - \eta(x)| \leq t) \leq C \cdot t^\alpha.$$

Note that  $\alpha = 0$  imposes no restriction on  $D$ , while larger values of  $\alpha$  impose greater restrictions. [9] showed the following result (adapted slightly to our setting, where the ranking risk is conditioned on  $Y \neq Y'$ ):

**Theorem 14** ([9]). *Let  $\alpha \in [0, 1]$  and  $q \in [1, \infty)$ . Then  $\exists$  a constant  $C_{\alpha, q} > 0$  such that for any distribution  $D$  on  $\mathcal{X} \times \{\pm 1\}$  satisfying noise assumption NA( $\alpha$ ) and any  $\hat{\eta} : \mathcal{X} \rightarrow [0, 1]$ ,*

$$\text{regret}_D^{\text{rank}}[\hat{\eta}] \leq \frac{C_{\alpha, q}}{p(1-p)} \left( \mathbf{E}_X[|\hat{\eta}(X) - \eta(X)|^q] \right)^{\frac{1+\alpha}{q+\alpha}}.$$

This allows us to obtain the following tighter version of our regret bound in terms of strongly proper losses under the same noise assumption:

**Theorem 15.** *Let  $\hat{\mathcal{Y}} \subseteq \mathbb{R}^*$  and  $\lambda > 0$ , and let  $\alpha \in [0, 1]$ . Let  $\ell : \{\pm 1\} \times \hat{\mathcal{Y}} \rightarrow \mathbb{R}_+^*$  be a  $\lambda$ -strongly proper composite loss. Then  $\exists$  a constant  $C_\alpha > 0$  such that for any distribution  $D$  on  $\mathcal{X} \times \{\pm 1\}$  satisfying noise assumption NA( $\alpha$ ) and any  $f : \mathcal{X} \rightarrow \hat{\mathcal{Y}}$ ,*

$$\text{regret}_D^{\text{rank}}[f] \leq \frac{C_\alpha}{p(1-p)} \left( \frac{2}{\lambda} \right)^{\frac{1+\alpha}{2+\alpha}} \left( \text{regret}_D^\ell[f] \right)^{\frac{1+\alpha}{2+\alpha}}.$$

*Proof.* Let  $c : \{\pm 1\} \times [0, 1] \rightarrow \mathbb{R}_+^*$  be a  $\lambda$ -strongly proper loss and  $\psi : [0, 1] \rightarrow \hat{\mathcal{Y}}$  be a (strictly increasing) link function such that  $\ell(y, \hat{y}) = c(y, \psi^{-1}(\hat{y}))$  for all  $y \in \{\pm 1\}, \hat{y} \in \hat{\mathcal{Y}}$ . Let  $D$  be a distribution



on  $\mathcal{X} \times \{\pm 1\}$  satisfying noise assumption  $\text{NA}(\alpha)$  and let  $f : \mathcal{X} \rightarrow \hat{\mathcal{Y}}$ . Then we have,

$$\begin{aligned}
\text{regret}_D^{\text{rank}}[f] &= \text{regret}_D^{\text{rank}}[\psi^{-1} \circ f], \quad \text{since } \psi \text{ is strictly increasing} \\
&\leq \frac{C_{\alpha,2}}{p(1-p)} \left( \mathbf{E}_X \left[ (\psi^{-1}(f(X)) - \eta(X))^2 \right] \right)^{\frac{1+\alpha}{2+\alpha}}, \\
&\quad \text{by Theorem 14, taking } q = 2 \\
&\leq \frac{C_{\alpha,2}}{p(1-p)} \left( \frac{2}{\lambda} \mathbf{E}_X [R_c(\eta(X), \psi^{-1}(f(X)))] \right)^{\frac{1+\alpha}{2+\alpha}}, \\
&\quad \text{since } c \text{ is } \lambda\text{-strongly proper} \\
&= \frac{C_{\alpha,2}}{p(1-p)} \left( \frac{2}{\lambda} \mathbf{E}_X [R_\ell(\eta(X), f(X))] \right)^{\frac{1+\alpha}{2+\alpha}} \\
&= \frac{C_{\alpha,2}}{p(1-p)} \left( \frac{2}{\lambda} \right)^{\frac{1+\alpha}{2+\alpha}} \left( \text{regret}_D^\ell[f] \right)^{\frac{1+\alpha}{2+\alpha}}.
\end{aligned}$$

The result follows by setting  $C_\alpha = C_{\alpha,2}$ . □

For  $\alpha = 0$ , as noted above, there is no restriction on  $D$ , and so the above result gives the same dependence on  $\text{regret}_D^\ell[f]$  as that obtained from Theorem 13. On the other hand, as  $\alpha$  approaches 1, the exponent of the  $\text{regret}_D^\ell[f]$  term in the above bound approaches  $\frac{2}{3}$ , which improves over the exponent of  $\frac{1}{2}$  in Theorem 13.

## 7 Conclusion and Open Questions

We have obtained upper bounds on the bipartite ranking regret of a scoring function in terms of the (non-pairwise) regret associated with a broad class of proper (composite) losses that we have termed *strongly proper* (composite) losses. This class includes several widely used losses such as exponential, logistic, squared and squared hinge losses as special cases.

The definition and characterization of strongly proper losses may be of interest in its own right, and may find applications elsewhere. An open question concerns the necessity of the regularity condition in the characterization of strong properness of a proper loss in terms of strong concavity of the conditional Bayes risk (Theorem 10). The characterization of strict properness of a proper loss in terms of strict concavity of the conditional Bayes risk (Theorem 4) does not require such an assumption, and one wonders whether it may be possible to remove the regularity assumption in the case of strong properness as well.

Many of the strongly proper composite losses that we have considered, such as the exponential, logistic, squared and spherical losses, are margin-based losses, which means the bipartite ranking regret can also be upper bounded in terms of the regret associated with pairwise versions of these losses via the reduction to pairwise classification (Section 3.1). A natural question that arises is whether it is possible to characterize conditions on the distribution under which algorithms based on one of the two approaches (minimizing a pairwise form of the loss as in RankBoost/pairwise logistic regression, or minimizing the standard loss as in AdaBoost/standard logistic regression) lead to faster convergence than those based on the other. We hope the tools and results established here may help in studying such questions in the future.

## Acknowledgments

Thanks to Harish Guruprasad and Arun Rajkumar for helpful discussions. Thanks also to Yoonkyung Lee for inviting me to give a talk at the ASA Conference on Statistical Learning and Data Mining held in Ann Arbor, Michigan, in June 2012; part of this work was done while preparing for that talk. This research was supported in part by a Ramanujan Fellowship from the Department of Science and Technology, Government of India.

## A Proof of Corollary 12

*Proof.* Let  $\hat{\eta} : \mathcal{X} \rightarrow [0, 1]$ . By Theorem 11, we have

$$\text{regret}_D^{\text{rank}}[\hat{\eta}] \leq \frac{1}{2p(1-p)} \mathbf{E}_{X, X'} \left[ |\eta(X) - \eta(X')| \cdot \mathbf{1}((\hat{\eta}(X) - \hat{\eta}(X'))(\eta(X) - \eta(X')) \leq 0) \right].$$

The result follows by observing that for any  $x, x' \in \mathcal{X}$ ,

$$(\hat{\eta}(x) - \hat{\eta}(x'))(\eta(x) - \eta(x')) \leq 0 \implies |\eta(x) - \eta(x')| \leq |\hat{\eta}(x) - \eta(x)| + |\hat{\eta}(x') - \eta(x')|.$$

To see this, note that the statement is trivially true if  $\eta(x) = \eta(x')$ . If  $\eta(x) > \eta(x')$ , then we have

$$\begin{aligned} (\hat{\eta}(x) - \hat{\eta}(x'))(\eta(x) - \eta(x')) \leq 0 &\implies \hat{\eta}(x) \leq \hat{\eta}(x') \\ &\implies \eta(x) - \eta(x') \leq (\eta(x) - \hat{\eta}(x)) + (\hat{\eta}(x') - \eta(x')) \\ &\implies \eta(x) - \eta(x') \leq |\eta(x) - \hat{\eta}(x)| + |\hat{\eta}(x') - \eta(x')| \\ &\implies |\eta(x) - \eta(x')| \leq |\hat{\eta}(x) - \eta(x)| + |\hat{\eta}(x') - \eta(x')|. \end{aligned}$$

The case  $\eta(x) < \eta(x')$  can be proved similarly. Thus we have

$$\begin{aligned} \text{regret}_D^{\text{rank}}[\hat{\eta}] &\leq \frac{1}{2p(1-p)} \mathbf{E}_{X, X'} \left[ |\hat{\eta}(X) - \eta(X)| + |\hat{\eta}(X') - \eta(X')| \right] \\ &= \frac{1}{p(1-p)} \mathbf{E}_X \left[ |\hat{\eta}(X) - \eta(X)| \right]. \end{aligned}$$

□

## References

- [1] S. Agarwal, T. Graepel, R. Herbrich, S. Har-Peled, and D. Roth. Generalization bounds for the area under the ROC curve. *Journal of Machine Learning Research*, 6:393–425, 2005.
- [2] N. Ailon and M. Mohri. An efficient reduction of ranking to classification. In *Proceedings of the 21st Annual Conference on Learning Theory*, 2008.
- [3] M.-F. Balcan, N. Bansal, A. Beygelzimer, D. Coppersmith, J. Langford, and G. B. Sorkin. Robust reductions from ranking to classification. *Machine Learning*, 72:139–153, 2008.
- [4] P. Bartlett, M. Jordan, and J. McAuliffe. Convexity, classification and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.

- [5] D. Buffoni, C. Calauzenes, P. Gallinari, and N. Usunier. Learning scoring functions with order-preserving losses and standardized supervision. In *Proceedings of the 28th International Conference on Machine Learning*, 2011.
- [6] A. Buja, W. Stuetzle, and Y. Shen. Loss functions for binary class probability estimation: Structure and applications. Technical report, University of Pennsylvania, November 2005.
- [7] C. J. C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton, and G. Hullender. Learning to rank using gradient descent. In *Proceedings of the 22nd International Conference on Machine Learning*, 2005.
- [8] S. Cl  men  on, G. Lugosi, and N. Vayatis. Ranking and empirical minimization of U-statistics. *Annals of Statistics*, 36:844–874, 2008.
- [9] S. Cl  men  on and S. Robbiano. Minimax learning rates for bipartite ranking and plug-in rules. In *Proceedings of the 28th International Conference on Machine Learning*, 2011.
- [10] S. Cl  men  on and N. Vayatis. Ranking the best instances. *Journal of Machine Learning Research*, 8:2671–2699, 2007.
- [11] C. Cortes and M. Mohri. AUC optimization vs. error rate minimization. In S. Thrun, L. Saul, and B. Sch  lkopf, editors, *Advances in Neural Information Processing Systems 16*. MIT Press, Cambridge, MA, 2004.
- [12] D. Cossock and T. Zhang. Statistical analysis of Bayes optimal subset ranking. *IEEE Transactions on Information Theory*, 54(11):5140–5154, 2008.
- [13] J. Duchi, L. Mackey, and M. I. Jordan. On the consistency of ranking algorithms. In *Proceedings of the 27th International Conference on Machine Learning*, 2010.
- [14] Y. Freund, R. Iyer, R. E. Schapire, and Y. Singer. An efficient boosting algorithm for combining preferences. *Journal of Machine Learning Research*, 4:933–969, 2003.
- [15] T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- [16] A. D. Hendrickson and R. J. Buehler. Proper scores for probability forecasters. *The Annals of Mathematical Statistics*, 42:1916–1921, 1971.
- [17] R. Herbrich, T. Graepel, and K. Obermayer. Large margin rank boundaries for ordinal regression. *Advances in Large Margin Classifiers*, pages 115–132, 2000.
- [18] T. Joachims. Optimizing search engines using clickthrough data. In *Proceedings of the 8th ACM Conference on Knowledge Discovery and Data Mining*, 2002.
- [19] W. Kotlowski, K. Dembczynski, and E. Huellermeier. Bipartite ranking through minimization of univariate loss. In *Proceedings of the 28th International Conference on Machine Learning*, 2011.
- [20] A. Rakotomamonjy. Optimizing area under ROC curves with SVMs. In *Proceedings of the ECAI-2004 Workshop on ROC Analysis in AI*, 2004.

- [21] P. Ravikumar, A. Tewari, and E. Yang. On NDCG consistency of listwise ranking methods. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics, 2010*, volume 15 of *JMLR Workshop and Conference Proceedings*, pages 618–626, 2011.
- [22] M. D. Reid and R. C. Williamson. Composite binary losses. *Journal of Machine Learning Research*, 11:2387–2422, 2010.
- [23] M. D. Reid and R. C. Williamson. Information, divergence and risk for binary experiments. *Journal of Machine Learning Research*, 12:731–817, 2011.
- [24] C. Rudin and R. E. Schapire. Margin-based ranking and an equivalence between adaboost and rankboost. *Journal of Machine Learning Research*, 10:2193–2232, 2009.
- [25] L. J. Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.
- [26] M. J. Schervish. A general method for comparing probability assessors. *The Annals of Statistics*, 17:1856–1879, 1989.
- [27] O. Stein. Twice differentiable characterizations of convexity notions for functions on full dimensional convex sets. IOR Preprint 2/2011, Institute of Operations Research, Karlsruhe Institute of Technology, 2011.
- [28] A. B. Tsybakov. Optimal aggregation of classifiers in statistical learning. *The Annals of Statistics*, 32(1):135–166, 2004.
- [29] K. Uematsu and Y. Lee. On theoretically optimal ranking functions in bipartite ranking. Technical Report 863, Department of Statistics, The Ohio State University, December 2011.
- [30] F. Xia, T.-Y. Liu, J. Wang, W. Zhang, and H. Li. Listwise approach to learning to rank: Theory and algorithm. In *Proceedings of the 25th International Conference on Machine Learning*, 2008.
- [31] T. Zhang. Statistical behavior and consistency of classification methods based on convex risk minimization. *The Annals of Statistics*, 32(1):56–134, 2004.